

1 REMOTE DOCUMENT UPDATING SYSTEM USING XML AND DOM

2 FIELD OF THE INVENTION

3 This invention relates to a system for updating and synchronizing a
4 document on a network server from a remote device. More specifically, this
5 invention uses the Document Object Model (DOM) specification to manipulate
6 documents, including databases that conform to the XML document structure
7 specification, to enable remote workstations, or clients, to update a database on a
8 server. Because the invention performs the update through the transmission of the
9 minimum amount of information necessary to fully update the server's database, it
10 is suitable for applications, including wireless transmissions, transmissions in
11 which the connection between the client and server is of limited bandwidth,
12 transmissions using conventional telephone lines, and transmissions through
13 computer networks using physical media.

14 BACKGROUND OF THE INVENTION

15 In modern business and industrial applications, it is imperative that critical
16 databases be updated and synchronized frequently and regularly to ensure the
17 accuracy of their information. It is common for entire offices and buildings to be
18 devoted to personnel whose primary task is to enter data into a database in the form
19 of additions of new data, deletions of old data, or modifications of existing data.

1 The data is entered from workstations ("clients") that are networked with a server
2 upon which the main database resides. Each change to the data is transmitted to
3 the server, which will temporarily lock the record from access by other users, make
4 and save the change, and then unlock the record so that it may be accessed by other
5 users. This process is repeated countless times during the course of a normal
6 workday, each change requiring essentially the same steps to be taken. When the
7 server and its clients are directly connected over a large bandwidth connection, the
8 overhead that may be inherent in multiple changes being made to the same record
9 may not be detrimental to the overall throughput of the system. However, when
10 the server and its clients must communicate over narrow bandwidth channels, the
11 possibility for congestion to detrimentally affect system performance is
12 significantly greater. The degree of congestion rises with the size of the database,
13 the number of entries needed to keep it current, and the frequency with which
14 accesses to the database are made.

15 The advent and increasing prevalence of wireless communications has made
16 it possible for databases to be maintained and brought current from remote sites.
17 Using wireless technology, critical databases can be maintained in a near-current
18 state from cellular telephones, hand-held devices, and personal data assistants
19 (PDA's), regardless of where they may be located. Common uses for such

1 technology may be found in the maintenance of personal and business calendars
2 from PDA's that are remote from the server where the organization's entire
3 calendar is maintained, sales automation and inventory management, and similar
4 applications. Uses for such technology could involve a customer's use of a
5 handheld device to register purchases in a supermarket; automated closeout of a car
6 rental upon return of the vehicle without the need for human oversight; or position
7 tracking for vehicles and aircraft based, not upon a radar/transponder reflection,
8 but upon the vehicle's own report of its position from global satellite positioning
9 data. Virtually every database application in which timely and accurate
10 information updating is critical is subject to remote data entry that may rely upon
11 wireless communications.

12 A primary drawback for such applications is that the bandwidth available for
13 database updates from remote clients may be limited. As wireless devices become
14 more prevalent, the availability of bandwidth for each application and user will
15 diminish, resulting in transmission bottlenecks, communications delays, and
16 ultimately, the untimely update of project-critical information. Since common
17 wireless applications such as cellular telephones and wireless Internet PDAs are
18 commonly assigned narrow bandwidths, there is a need for a system that will
19 reduce the number of accesses from client to server and that will require the

1 transmission of only the smallest amount of data that is necessary to provide
2 complete information for the server to update its database. The concept of sending
3 only the smallest amount of information needed to perform the update is known as
4 "granularity." In general, a finer granularity results in a more efficient the use of
5 the available bandwidth.

6 SUMMARY OF THE INVENTION

7 The present invention is a system that resides on the remote client, and
8 which notes and records the accumulation of mutations (additions, deletions,
9 modifications) to the client's local database, which normally will constitute a
10 subset of the main database that resides on the server. Before changes to the local
11 database are transmitted to the server, they are processed and reduced to the
12 smallest amount of data, or the finest granularity, that will allow the server to
13 record the changes to the main database. The system uses XML (eXtensible
14 Markup Language) documents to represent the local database, and further uses the
15 DOM (Document Object Model) established by the World Wide Web Consortium
16 (W3C), to provide a standardized interface for manipulation of the XML
17 document. Although the specifications for the DOM and the XML are currently
18 undergoing a process of development and refinement, the implementation of those
19 specifications in this invention is not intended to be version-specific, but is generic

1 and should perform effectively with future versions of those specifications as well
2 as with all existing versions.

3 In XML, a "document" may represent a collection of related data, including
4 a database. In the DOM specification, an XML document may be represented as a
5 logical "tree" having objects, or "nodes," located in a hierarchical branching
6 structure. Each node has various properties, including at least a "name" (or
7 "label") property, and a "data value" property. Each discrete data element in the
8 XML document can be identified by an XML tag, and the DOM will have a node
9 corresponding to each such data element. When data in the XML document is
10 added, deleted, or modified, the DOM also changes to reflect the modification.
11 The automatic modification of the DOM to conform to the XML document is a
12 property of the DOM specification and its implementation. Mutations to the XML
13 document are recorded in an Event Table as "events." All events are stored, cross
14 referenced to the path of the node, a timestamp, and the mutation type. When data
15 synchronization occurs, the Event Table is parsed, and each node at which multiple
16 events took place is collapsed to a single event using a set of rules described below.
17 The data value of such collapsed node is the value left by the last mutation
18 involving the node. By collapsing the number of node events in the DOM after
19 modifications have been made to the document, the data necessary to update the

1 main database is brought to a minimum, and may then be transmitted to the server
2 for application to the main database.

3 The process of collapsing the number of node events requires a sorting,
4 followed by the parsing of each event in the Event Table whereby events to be
5 passed to the remote database will be saved in a Save Table, while redundant
6 events are discarded. When the Save Table has been created, a synchronization
7 manager will then retrieve data from each node referenced in the Save Table, and
8 the data for each event in the Save Table will be transmitted to the main database
9 along with the node address, an optional date stamp, and the event type.

10 Accordingly, it is an object of the present invention to provide a system for
11 tracking changes made at a workstation to a subset of a database and, before
12 sending the changes to the server, to reduce to a minimum the amount of
13 information that must be sent. It is a further object of this invention to increase the
14 efficiency of database updates over narrow bandwidth communication channels. It
15 is another object of the invention to utilize the XML and DOM specifications to
16 process data on a remote client using a standardized format for reducing the
17 amount of information necessary to fully describe the changes to the database.
18 These and other objects of the invention will become evident through the following
19 description of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a flow chart depicting the flow of events through a remote client prior to transmission to a server for entry into a database.

FIG. 2a shows an example in which the local database is an address book.

FIG. 2b shows a logical representation of the address book in the format of an XML document.

FIG. 3a lists hypothetical changes that are made to the address book.

FIG. 3b shows the hypothetical changes to the address book as they are recorded in the Event Table.

FIG. 4 shows the Event Table after it has been sorted and before it has been parsed.

FIGs. 5a and 5b are connecting sections of a flow chart depicting one embodiment of the decision-making process that will reduce the number of events that will be sent to the main database. In FIGs. 5a and 5b, the Event Table has been sorted, but has not yet been parsed.

FIG. 6 shows the Save List for mutations made to the hypothetical address book of FIGs. 2 - 4.

FIG. 7 is a DOM representation of the XML document of FIG. 2b prior to changes having been made.

FIG. 8 is the DOM representation of the XML document representing the hypothetical address book after changes have been made to the address book.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

As is shown in the flow chart of FIG. 1, changes made to a local database 100 are registered as "events" by Event Listener 110. In the XML-DOM specification, as applied in the invention, an event occurs whenever a node is added or deleted, or the data in a field is modified. In FIG. 1, events 120 are detected at Event Listener 110 and passed to Change Manager 130 where they are listed as records in Event Table 300. When modifications to the XML database have been completed, the Change Manager 130 parses Event Table 300 to determine which nodes in the DOM tree are affected by the events listed in the Event Table. The parsing creates a second list of changes (the "Save Table") to be transmitted to the computer on which the main database resides which, when applied to the main database, will update that database with the remotely-performed changes, thus keeping the client and server data in synchronization. This second list represents the smallest number of modifications that must be made to fully update the database resident on the server. The Save Table may be

1 compressed with any commercial compression product for transmission to the
2 server where the database will be modified in accordance with data received from
3 the Synch Manager.

4 In FIG. 2a, an example is shown using an address book as the local database
5 200. It will be understood, however, that this invention is not limited to this
6 exemplary implementation, and that any document capable of being formatted as
7 an XML document having a DOM representation, and subject to multiple
8 modifications made from a remote terminal, will fall within the scope of the
9 invention. In this example, the address book may constitute a subset of a larger
10 organization-wide address book maintained on the server. The local address book
11 has three fields: "Name" 210, "Phone" 220, and "E-Mail" 230. Three records are
12 in the address book, and the manner in which the system of this invention updates
13 the organization-wide address book with changes made to the records shown in
14 FIG. 2a will be described.

15 FIG. 2b shows the address book represented as an XML document. In the
16 XML specification, a "document" may constitute almost any object having
17 properties that include a value. The XML structure reflects the organization of the
18 database. Thus, the XML tag <address book> 250, represents the database, and is
19 located at the highest level. Each database record 260, defined as a "person," is

1 located at an intermediate level, and has a unique ID attribute that uniquely
2 identifies it and distinguishes it from other elements at the same level. Discrete
3 data elements representing the data in each field of each record are located at the
4 lowest level. In FIG. 2b, the XML tag <person ID= "1"> 260 is at the intermediate
5 level, while data maintained under the XML tags for <name> 270, <phone> 280,
6 and <email> 290 are located at the lowest level.

7 The DOM presents documents as a hierarchy of “node” objects. The DOM
8 corresponding to the XML document of FIG. 2b is a branching hierarchy, as
9 depicted in FIG. 7. The DOM may include a variety of node objects that also
10 implement other, more specialized interfaces for accessing and manipulating
11 document objects. Some types of nodes may act as “parent” nodes having “child”
12 nodes below them in the hierarchy. Others may be “leaf” nodes that cannot have
13 anything below them in the document structure. Each node in the DOM
14 corresponds to an XML tag. Each node has properties, including at least a “name”
15 property and a data value property. As shown in FIG. 7, node 411 is an “element”
16 type object, and has a name property (“Person”) and an ID attribute 530 of “1.”
17 The node is a parent-type node established in the DOM specification. Nodes 414 -
18 422 are additional element nodes that do not have unique ID attributes associated
19 with them. However, for each sub branch of the DOM, the element name property

for each element will not be repeated in that sub branch. Nodes 422-431 are leaf “text” nodes and cannot have child nodes below them. A text node contains only a data value and is associated with an element node. For example, node 416 is the element node having the property “phone,” and node 425 is the associated text containing the data “617-321-7654.” Another example is shown at nodes 418 and 427 having the element “email,” and a data value of “ddoe@thatdomain.com.”

Mutations are made to the local database by adding, deleting, or substituting data elements in the XML document. Example changes to the address book are shown in tabular form in FIG. 3a, while FIG. 3b shows an Event Table that corresponds to those changes. The Event Table has fields that loosely correspond to the fields of the address book. Thus, for example, the “Event” column 340 of the address book, which identified objects undergoing a change, corresponds to the XQL path 310 to the node in the DOM at which the object that was the subject of the event is located; the “Change” column 350 of the address book describing the change corresponds to the “Nature of Event” field 320 (whether the event was a record addition or deletion, or a field modification) in the Event Table; and the “Time” column 360 describes the time the change was made to the address book, and corresponds to the “Time” field 330 in the Event Table. Although each instance of change in the database is recorded in the Event Table, specific data that

1 is added, deleted, or substituted is not so recorded. Rather, once the Event Table
2 has been processed and the number of events has been reduced to a minimum, the
3 Change Manager will retrieve the relevant data values from the DOM to pass to the
4 Sync Manager for transmission to the server. Each addition, deletion, or
5 modification of data in the address book will cause a new record to be added to the
6 event table until the synchronization process begins. Following synchronization,
7 the data in the Event Table is cleared and new data will be added as further
8 mutations are made to the local database.

9 When a new record is to be added to the local database, only such
10 information as is needed to update the main database will be saved and transmitted.
11 For example, in FIG. 3a, it will be seen that a new person, "Carol Roe" is being
12 added to the address book 372. In order to minimize storage space on the client,
13 the default method for adding a new record to the local database will not add either
14 element or text nodes that are currently not being used. Therefore, neither element
15 nor text nodes are added at event 372. A node that exists in the XML Document
16 Type Definition (DTD), but not in the DOM tree, represents a null node value.
17 Thus, when a new record is to be added to the address book, initially only a new
18 "person" element is added. Thereafter, when a name 373, phone 342, or email 375

1 is added to the address book, the corresponding child element and text nodes will
2 be added to the Event Table, 383, 384, and 385, and to the DOM.

3 The DOM of FIG. 7 depicts parent-child relationships between the nodes as
4 a static, solid connecting line. However, the DOM structure is a "live"
5 construction, meaning that it is constantly changing to reflect the contents of the
6 XML document in real time as the document is modified. In FIG. 7, the DOM
7 shows the address book as it exists in FIGs. 2a and 2b, before any of the changes
8 shown in FIGs. 3a and 3b have been made. Each node has a name property, a data
9 value property, and a type property. The type properties of the DOM describe the
10 type of node. The Worldwide Web Consortium (W3C) has defined the basic
11 element, text, attribute, and document node types. The name properties of the
12 DOM in this example have been selected to be descriptive of the nodes. In FIG. 7,
13 the DOM has a root node 410 of type "Document" having the name property
14 ("address book"), but has no data value. Three "Element" type nodes 411 - 413 are
15 situated immediately below the document node and have the name property
16 "person". Each person "element" node has an attribute node of name "ID" and a
17 value to uniquely distinguish it from every other person "element" node. Each
18 person "element" node has three child "element" nodes, corresponding to the data
19 in each of the database fields "name," "phone," and "email" for the relevant record.

1 Text nodes 423 through 431 in FIG. 7 represent the data values for the parent
2 “element” nodes.

3 In accordance with the DOM specification, nodes may be interfaced and
4 manipulated with the methods that are applicable to each type of node. Although
5 the DOM specification defines many such node types, the only nodes represented
6 in FIG. 7 correspond to document, element, attribute, and text nodes. Returning to
7 FIG. 1, changes made to the local database are implemented as changes to the
8 XML document 100 representing the address book, and trigger events 120 that are
9 detected by the Event Listener 110 and passed on to the Change Manager 130. The
10 Change Manager creates and manages the Event Table 300 in which all events
11 occurring after the most recent synchronization are listed. The address book
12 shown in FIG. 3a and the Event Table shown in FIG. 3b each list eight events
13 reflecting changes in the local database. Changes to the address book are identified
14 at 370 - 377, while corresponding records entered into the Event Table are
15 identified at 380 - 387. At 370 it may be seen that David Doe changed his phone
16 number. That event is reflected in the Event Table as an XQL path identifying the
17 "phone" node for the person having ID = "2" 380 and a “change” of data. Next in
18 the list, Sally Jones changed her e-mail address 371, which corresponds to event
19 381 in the Event Table. A new, null record 372 was added to the address book to

1 create a record for Carol Roe, and is shown in the Event Table at 382. Carol's
2 name 373, phone 374, and e-mail 375 are added in the address book, and
3 correspond to events 383, 384 and 385 in the event table of FIG. 3b. Next, David
4 Roe is removed from the address book 376, and this event is shown as a deletion
5 386 in FIG. 3b. Lastly, Carol Roe changed her phone number 377, and this
6 modification is shown in the Event Table at 387. This is the last event before the
7 changes will be processed, and the database on the server will be brought into
8 synchronization with the modified data in the client address book.

9 The Change Manager determines which data must be passed to the server
10 through a process of sorting and parsing the Event Table, and referring to the
11 modified DOM to obtain the data related to each event in the Save Table. This
12 process may follow any one of a number of methods for obtaining relevant data,
13 and the steps outlined here are exemplary only, and do not represent the only
14 means for obtaining such data. Regardless which method the Change Manager
15 applies, however, it will produce the minimal change set needed to bring the server
16 database into synchronization with that of the client.

17 In the embodiment depicted in FIGs. 2 and 3, the Change Manager will first
18 sort, and then parse the Event Table to generate the minimal amount of data needed
19 to send to the server. One method for doing this is to sort the Event Table by the

1 XQL (structured query language for XML) path and timestamp. When sorted in
2 alphabetical and timestamp order, the Sorted Event Table will place all events
3 related to the same node consecutively in the table, and will place the earlier
4 occurring events first. In this fashion, a node that has been added and later
5 modified will appear in the table with the ADD event being above any subsequent
6 MODIFY or DELETE events pertaining to the same node. A modified XQL is
7 sorted alphabetically to group together all child nodes and branches, regardless of
8 their location in the Event Table. FIG. 4 shows the address book example of FIG.
9 3b after it has been sorted alphabetically and by timestamp.

10 All events in the Event Table 300 begin with the XQL path prefix
11 "//person[@ID=", such that the remainder of the XQL path will determine the
12 listing order. As is shown in FIG. 4, person element nodes having an "ID" value of
13 "2" are grouped together (380 with 386) followed by the sole person element node
14 having an "ID" value of "3" (381). The largest grouping is of person element
15 nodes having an "ID" of "4" (382, 383, 384, 385, and 387). Within each grouping,
16 the shortest XQL path will have sorted to the top of the group, while progressively
17 longer XQL paths will appear lower in the group listing. In order for the
18 alphabetical sorting to work, the XQL path may be slightly modified by adding a
19 special character to the beginning of the element name, which will cause the

1 element to move to the bottom of the alphabetical list when sorted. After sorting,
2 the shortest XQL path will correspond to the highest parent node of a branch, as
3 with nodes at 382 and 386, and will appear at the top of a grouping in the table.
4 Next will be individual elements that do not have child elements, such as nodes
5 383 - 385. Last will be nodes that have further child nodes, although none are
6 depicted in the example of an address book. A deep node will have a long XQL
7 path, but the path includes information about its parent nodes and will be sorted
8 appropriately. After being sorted, the Event Table of FIG. 3b will have all of the
9 events in order, as shown in FIG. 4. The Change Manager will then parse the list
10 to determine which events should be entered in the Save Table.

11 As depicted in FIGs. 5a and 5b, the Change Manager follows a set of rules to
12 determine the minimal set of data to send to the server. Other embodiments may
13 follow the same general steps and fall within the scope of this invention, yet differ
14 in the specific manner in which they accomplish the steps. While other processing
15 methods and steps may be used, one specific embodiment of the process is set forth
16 in FIGs. 5a and 5b.

17 The general process for generating the Save Table involves a comparison of
18 two events, denoted as Event A and Event B, which are designated as consecutive
19 events in the sorted Event Table. Certain cases, however, require that save

1 decisions be based upon the analysis of a single event, such as when the Event
2 Table includes only a single event, or when there is but a single event remaining to
3 be processed. The condition in which there is only one event in the Event Table
4 will be detected at 600, where an End of File flag will cause processing to follow
5 steps shown at 610 - 650. If the single event is an ADD 610, then that event will
6 be entered in the Save Table 630 as an ADD, and processing will end 650. If the
7 single event is a MODIFY 620, then that event will be entered in the Save Table
8 635 as a MODIFY, and processing will end 650. If the single event is a DELETE,
9 then the event will be entered as a DELETE event 640, and processing will end
10 650.

11 Where the Event Table includes two or more events, processing will proceed
12 to 660, where the next event (a B event) will be compared with event A. If A and
13 B have the same path 670 then A can only be an ADD or a MODIFY, while B can
14 only be a MODIFY or a DELETE. If B is a MODIFY 680, then A is kept and B is
15 skipped or discarded 700. If B is a DELETE, then the "Delete Flag" (DF) is set to
16 "True" 690, and A is kept while B is discarded 700. Next, the EOF flag 710 is
17 checked and, if the end of file has been reached, the process for handling a sole
18 remaining event will be followed, as is described below. If there is another event

1 in the Event Table, a new event will be retrieved as a B event 730, and the
2 comparison process will recommence 660.

3 If A and B do not have the same path 670, they will next be checked to see
4 whether B is a subset of A 740, that is, whether B falls along a path that is a sub
5 branch of the path upon which A is located. If B is a subset of A, the process next
6 determines whether A is an ADD 750. If it is, the process determines whether B is
7 also an ADD 760. If both A and B are ADDs, the process next determines whether
8 the DF (delete flag) is True 770. If the DF is True (DF is set), then A is kept and B
9 is discarded 700, and the process checks for another event 710. Alternatively, if
10 the DF is False 770, then A is saved as an ADD and B becomes the next A 790,
11 and the EOF 710 is checked.

12 If A is an ADD and B is not an ADD 760, then B must be a MODIFY. The
13 reason that B is a MODIFY, and cannot be a DELETE is that a child “element”
14 node representing a name, an e-mail, or a phone, cannot be deleted independently,
15 but can be deleted only as part of the deletion of its parent “person” element node.
16 While the deletion of the parent node is recorded as an event, the deletion of child
17 nodes below the parent are not written as events in the Event Table. Having
18 determined that B is a MODIFY 760, the process will next check the DF 780. If
19 the DF is True (set), A is kept 700, B is discarded, and a next event is looked for

1 710. If the DF 780 is False (not set), then A is added to the Save Table as an ADD
2 790 and event B becomes event A for purposes of the next comparison.

3 A similar process is followed if B is a subset of A 740, and A is not an ADD
4 750, but is a MODIFY 800. If B is an ADD 810, then the DF flag is checked 820
5 and, if True, A is kept 700 and B is discarded. If the DF is False 820, then A will
6 be saved as a MODIFY 795 and B will become the next A. If B is not an ADD
7 810, then it must be a MODIFY, and the DF is checked 830. If the DF is True 830,
8 then A will be kept 700 and B will be discarded. If DF is False, then A will be
9 saved as a MODIFY and B will become the next A 795.

10 If B is a subset of A 740, and A is not an ADD 750 or a MODIFY 800, then
11 A must be a DELETE. Because of the manner in which the Event Table is sorted,
12 it is possible that A may be a DELETE while B is an ADD or a MODIFY.

13 Regardless whether B is an ADD or a MODIFY, the processing will result in A
14 being kept 700, and B being discarded. This result flows from the fact that, if A is
15 DELETED, any nodes below it in the same branch will necessarily also be deleted
16 and will not require processing.

17 When all but the last event in the Event Table have been processed, an EOF
18 flag 710 will signify that there are no other events to be retrieved. There will be a
19 single remaining A event that must then be processed. If that event is an ADD and

1 the DF is False 900, then it will be saved as an ADD 630. If the A is an ADD and
2 the DF is True 900, then A will be saved as a DELETE. If A is not an ADD, it is
3 then checked to see whether it is a MODIFY 890. If it is a MODIFY and the DF is
4 False 910, then A will be saved as a MODIFY. If the DF is True 910, then A will
5 be saved as a DELETE. If A is neither an ADD 880 nor a MODIFY 890, then it
6 will be saved as a DELETE 640. Once the event has been saved, processing of the
7 Save Table terminates 650 and the Synch Manager 140 will retrieve data values
8 from the DOM corresponding to ADD and MODIFY events in the Save Table.

9 Turning next to FIG. 5b, processing of events is shown where the XQL paths
10 for A and B are not within the same sub branch. In general, where the paths are
11 different, the process will Save event A and will move event B to A and retrieve
12 the next event as B. Because the B event represents a path that is different from
13 that of A, the result of processing will leave the DF flag clear (False), will save the
14 A event, and will move the B to the A event.

15 Thus, if A is an ADD 910, then DF 920 is checked and, if True, the DF flag
16 is cleared (set to False) 930 and A is saved as a DELETE 940 and event B becomes
17 event A. Processing then proceeds to check for an EOF condition 710, and further
18 processing takes place as shown following EOF 710 in FIG. 5a. If DF 920 is
19 False, A will be saved as an ADD 950 and event B becomes event A.

1 If A is a MODIFY 960, then DF 970 will be checked and if it is False, A will
2 be saved as a MODIFY 980. If DF 970 is True, A will be saved as a DELETE
3 940. In either case, event B then becomes the next event A.

4 If A is not a MODIFY 960, then it must be a DELETE and A will be
5 saved as a DELETE while B becomes the next A 940. The DF will be cleared 930,
6 and processing will continue with a check for the EOF 710.

7 These decision rules shown in FIGs. 5a and 5b may be followed through the
8 Sorted Event Table of FIG 4 to result in the Save List shown in FIG. 6. The
9 Change Manager will retrieve the first event [386], which is a DELETE event, will
10 designate it as a A event, and will check to ensure that there is more than one event
11 600. The Change Manager will next retrieve 660 the B event [380], which is a
12 MODIFY event [380]. At step 670, the Change Manager determines that the
13 events do not have the same path 670. Proceeding to step 740, the Change
14 Manager determines that B is a subset of A 740, then determines that A is not an
15 ADD 750 or a MODIFY 800. The Change Manager then keeps (but does not yet
16 save) A as a DELETE 700, discards B, and retrieves 660 the next B event [381].

17 B event [381] is a MODIFY that is not a subset of A. The Change Manager
18 proceeds to determine that A is not an ADD 910 and is not a MODIFY 960. The
19 Change Manager then clears the DF 930, saves A as a DELETE, and moves B to A

1 940. It then retrieves a new B event 660, which is an ADD event [382] that is not a
2 subset of A. Proceeding to steps 910 and 960, the Change Manager determines
3 that new event A [381] is a MODIFY 960. The DF is False 970, A is saved as a
4 MODIFY and B becomes new event A 980.

5 The next B event [385] is an ADD that is determined to be a subset of event
6 A [382] at step 740. Both A and B are ADDs 750 and 760, and the Change
7 Manager next checks DF 770. Since DF 770 is clear, event A is saved as an ADD
8 and event B moves to be new event A.

9 The next B event [383] is an ADD that is on a different sub branch from
10 event A[385]. The Change Manager will process event A at steps 910 and 920,
11 and will save event A as an ADD 950. Event B becomes new event A 950, and the
12 next event in the Sorted Event Table [384] is retrieved. This, too, is an ADD that
13 is on a different sub branch, and processing at steps 910 and 920 will cause event
14 A to be saved as an ADD and event B to become new event A 950.

15 The change manager will then retrieve the last event [387] in the Sorted
16 Event Table as a B event, and will determine that it has the same path as event A
17 670. Since A is an ADD and B is a MODIFY 680, A will be kept and B will be
18 discarded 700. Upon checking EOF 710, the Change Manager will determine that
19 no further events remain to be retrieved, and processing will continue at step 880.

1 Since event A is an ADD, and the DF is not set 900, A will be saved as an ADD
2 630, and processing of the Sorted Event List will terminate 650. At that time, the
3 Save List is complete and ready for transmission to the server.

4 FIG. 6 shows the Send List after it has been alphabetically sorted and then
5 processed by the Change Manager. Comparing the list of FIG. 6 with that of FIG.
6 4, it may be seen that two changes made to the XML document need not be sent to
7 the database on the server, and that the Send List represents the least amount of
8 data that will synchronize the database on the server to that on the client device.

9 FIG.s 7 and 8, respectively, provide a depiction of the DOM before, and then
10 after modifications were made to the address book. In FIG. 7, the DOM shows
11 three elements 411, 412, and 413, each of which is identified by a tag ("ID") 520,
12 530, and 540, having a different value. Each of the three records has three data
13 elements: Name, 414, 417, and 420; e-mail, 415, 418, and 421; and phone, 416,
14 419, and 422. The values of the data elements are shown, 423 – 431, and the DOM
15 tree of FIG. 7 provides a complete picture of the address book.

16 In FIG. 8, modifications have been made to the address book, and the DOM,
17 which is a "living" object that reflects changes as they are made in real time, has
18 been modified to reflect those modifications. In FIG. 8, the element 412 having an
19 "ID" = "2" 520 has been deleted, as is symbolized by the dashed line separating

1 that element from the DOM. Although an earlier modification was made to the
2 telephone number for that element, the modification was later subsumed by the
3 deletion of the entire record.

4 In addition, modifications made to the element having an "ID" of "3" are
5 seen at 430, in a modification of the e-mail address for that element. Lastly, the
6 addition of a new element 510 having an "ID" of "4" 550 is shown in the DOM of
7 Figure 8. The data values existing in the DOM at the time of synchronization will
8 all be obtained at one time and sent to the server, along with the Save Table for
9 inclusion in the database residing on that server.

10 A further aspect of the invention may require the server to send updated
11 information from its database to the client's local database for processing. This
12 could occur, for example, when a single master database receives updates from
13 more than one source. The processes of this invention may be employed at the
14 server to reduce the amount of data that must be transmitted before the
15 transmission to the client is performed. Potential conflicts that may occur when the
16 same record is updated from two remote sites prior to synchronization may be
17 handled by methods and processes known in the art.

18 It will be understood that the invention is not limited to the processes shown
19 and described, which are exemplary, and that other uses and configurations may be

1 employed that will fall within the spirit and scope of the invention. For example,
2 although the processing steps have been described with reference to an
3 alphabetized and time sorted event list, an alternative embodiment would be to sort
4 each group of identical XQL paths to place all DELETE events ahead of all ADD
5 or MODIFY events in the list, thereby eliminating some of the steps required to
6 process the sorted list. Similarly, other schemes for sorting the list or reducing the
7 amount of data that must be transmitted to the server to a minimum may be
8 employed, which may be considered to fall within the scope of this invention.

9 While the invention has been described, disclosed, illustrated and shown in
10 various terms or certain embodiments or modifications which it has assumed in
11 practice, the scope of the invention is not intended to be, nor should it be deemed
12 to be, limited thereby and such other modifications or embodiments as may be
13 suggested by the teachings herein are particularly reserved especially as they fall
14 within the breadth and scope of the claims here appended.